

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
9 January 2003 (09.01.2003)

PCT

(10) International Publication Number
WO 03/003236 A1

(51) International Patent Classification⁷: G06F 15/173,
17/00

(21) International Application Number: PCT/IB02/02507

(22) International Filing Date: 27 June 2002 (27.06.2002)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
09/895,532 29 June 2001 (29.06.2001) US

(71) Applicant: NOKIA, INC. [US/US]; 6000 Connection
Drive, Irving, TX 75039 (US).

(72) Inventor: WANG, Bing; 233 Blossom Hill Road, Apart-
ment 6, San Jose, CA 95123 (US).

(74) Agent: WRIGHT, Bradley, C.; Banner & Witcoff, Ltd.,
11th floor, 1001 G Street, N.W., Washington, DC 20001-
4597 (US).

(81) Designated States (national): AE, AG, AL, AM, AT, AU,
AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU,

CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH,
GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC,
LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW,
MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SD, SE, SG,
SI, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, UZ, VN,
YU, ZA, ZM, ZW.


(84) Designated States (regional): ARIPO patent (GH, GM,
KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW),
Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM),
European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR,
GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent
(BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR,
NE, SN, TD, TG).

Published:

- with international search report
- before the expiration of the time limit for amending the
claims and to be republished in the event of receipt of
amendments

For two-letter codes and other abbreviations, refer to the "Guid-
ance Notes on Codes and Abbreviations" appearing at the begin-
ning of each regular issue of the PCT Gazette.

(54) Title: APPARATUS AND METHOD FOR HANDLING ELECTRONIC MAIL



(57) Abstract: A mail transfer agent adapted for determining the disposition of incoming e-mail from a sender includes a penalty count filter module (20) which functions on the basis of current communication system resource usage, such as the number of concurrent TCP connections being maintained, and penalty counts assigned for cumulative undesirable sender activity such as sending large e-mail files. The penalty count filter module can operate in a selective-rejection state accepting e-mail from senders without a penalty count and randomly accepting other e-mail, and in a random-rejection state rejecting all e-mail from senders with a penalty count and randomly rejecting other e-mail. The rejection rates can be increased or decreased in response to a detected increase or decrease in the usage of system resources.

WO 03/003236 A1

Title: Apparatus and Method for Handling Electronic Mail

FIELD OF THE INVENTION

This invention relates to network communication systems and, in particular, to a system and method for handling incoming electronic mail messages.

BACKGROUND OF THE INVENTION

Denial of service attacks, including mail flooding, are common problems affecting the security of a mail transfer agent used in the management of electronic mail (e-mail). In the present state of the art, one possible response is to reject all incoming e-mail during a mail flooding or denial of service attack incident. However, such a response may be undesirable as e-mail from a legitimate sender can be rejected along with the mail flood from a problem sender. The present state of the art does not provide a method of distinguishing between the problem sender and the legitimate sender during periods of high e-mail activity.

What is needed is a method for responding to incoming e-mail which selectively rejects e-mail from a problem sender and selectively accepts e-mail from a legitimate sender during the onset of mail flooding or denial of service attack.

SUMMARY OF THE INVENTION

The present invention results from the observation that e-mail senders initiating undesirable activities can be identified and tracked by utilizing a penalty count filter module, integrated into the front end of a mail transfer agent in a communication device, so that subsequent e-mail disruptions to the receiving device can be mitigated. The penalty count filter module determines the disposition of incoming e-mail on the basis of current communication system resource usage and penalty counts assigned for cumulative undesirable sender activity. System resource usage may be determined by the number of concurrent TCP connections being maintained, and the undesirable sender activity, which is tracked using a behavior

trace table, may include sending a large number of e-mails or using a relatively large amount of TCP connection time. The penalty count filter module operates in a plurality of states, including a normal state, a selective-rejection state, and a random-rejection state. In the selective rejection state, e-mail from senders having a penalty
5 count are randomly rejected, and e-mail from senders without a penalty count are accepted. In the random-rejection state, all e-mail from senders having a penalty count is rejected, and e-mail from senders without a penalty count are randomly rejected. The rejection rates can be increased or decreased in response to a detected increase or decrease in the usage of system resources.

BRIEF DESCRIPTION OF THE DRAWINGS

10 The invention description below refers to the accompanying drawings, of which:

Fig. 1 is a diagrammatical representation of a mail transfer agent including a penalty count filter module for determining the disposition of incoming electronic mail;

15 Fig. 2 is a flow diagram describing the sequence of operations performed by the penalty count filter module of Fig. 1;

Fig. 3 is a data structure diagram of a penalty count table resident in the penalty count filter module of Fig. 1;

20 Fig. 4 is a flow diagram describing in greater detail the operation of the penalty count filter module, as represented by the step of determining filter module status in the flow diagram of Fig. 2;

Fig. 5 is a flow diagram describing in greater detail the operation of the penalty count filter module, as represented by the step of processing e-mail in accordance with sender and filter module status in the flow diagram of Fig. 2;

Fig. 6 is a diagrammatical illustration of a communications network including the penalty count rejection table of Fig. 3 resident in a mail transfer agent and a behavior trace table resident in a workstation;

Fig. 7 is a data structure diagram of the behavior trace table of Fig. 6;

5 Fig. 8 is a data structure diagram of a merged penalty count table which includes data from the penalty count table of Fig. 3 and the behavior trace table of Fig. 7; and

Fig. 9 is a flow diagram describing an alternative sequence of operations performed by the penalty count filter module of Fig. 1.

DETAILED DESCRIPTION OF AN ILLUSTRATIVE EMBODIMENT

10 There is shown in Fig. 1 a functional block diagram of a mail transfer agent 10, adapted for receiving an incoming e-mail message 11 via a communications network (not shown) such as the Internet. The mail transfer agent 10 can be integrated with a communication appliance, such as a personal computer or a workstation, for example. The mail transfer agent 10 includes a penalty count filter
15 module (PCFM) 20 integrated into a front end 15 of the mail transfer agent 10, an incoming message queue 13, and a forwarding daemon 18. In a preferred embodiment, the front end 15 functions in accordance with Simple Mail Transfer Protocol (SMTP) as is well known in the relevant art.

The incoming e-mail message 11 is provided to an accept/reject filter 23
20 which determines whether to save the incoming e-mail message 11 to the incoming message queue 13 as an accepted e-mail message 16, or reject the incoming e-mail message 11 and transmit a transient negative completion reply 19 to the originator of the message. The accept/reject determination is made by the accept/reject filter 23 on the basis of a sender penalty count status 27 and a PCFM state 29. As explained in

greater detail below, a sender identifier 21 is obtained from the incoming e-mail message 11 and is used to determine the sender penalty count status 27.

The sender may be identified by using a reverse Domain Name Service (DNS) verification to ascertain the IP address of the originator of the incoming e-mail message 11. Under circumstances in which the envelope address of the sender is not available, such as during a TCP timeout attack, the peer IP address of the TCP connection can be used as the sender identifier 21. The sender penalty count status module 27 maintains a constantly-updated list of senders associated with undesirable e-mail activity. Such undesirable activity may include, for example, sending large numbers of e-mails, sending e-mails of relatively large sizes, using too much TCP connection time, or causing a TCP timeout.

The system resource usage status 25, which is updated whenever a TCP connection is established, is used to determine the PCFM state 29, as described in greater detail below. The system resource usage status 25 provides a value for the usage or capacity status of one or more system resources related to the processing of incoming messages, including the incoming e-mail message 11, such as disk space occupied by the incoming message queue 13, the number of e-mail files in the incoming message queue 13, or the number of concurrent TCP connections being maintained.

Operation of the mail transfer agent 10 can be described with additional reference to the flow diagram of Fig. 2, in which the mail transfer agent 10 completes a TCP connection, at step 31, and receives the incoming e-mail message 11. The PCFM state 29 is determined, at step 33, and the identity of the sender of the e-mail message 11 is determined, at step 35. The PCFM state 29 is determined as a function of: i) the current PCFM state, ii) the time period for which the penalty count filter module 20 has remained in the current PCFM state, and iii) the current system

resource usage status 25. Determination of the PCFM state 29 is not dependent on the identity of the sender of the e-mail message 11.

Using a process described in greater detail below, the accept/reject filter 23 determines, at step 37, either that the incoming e-mail message 11 is to be transmitted to the addressee as a forwarded e-mail message 17, or that the incoming e-mail message 11 is to be rejected and an optional transient negative completion reply 19 is to be transmitted to the sender, stating that the incoming e-mail message 11 has been rejected. The determination of accepting or rejecting the e-mail message 11 is made based on the current PCFM state 29 and the identity of the sender of the e-mail message 11. When the current incoming e-mail message 11 has been processed, the mail transfer agent 10 updates a behavior trace table 160 (shown in Fig. 7, below), at step 41. In a preferred embodiment, the activity of each e-mail sender is tracked and updated by means of the behavior trace table 160. It should be understood that not all e-mail senders listed in the behavior trace table 160 will subsequently produce sufficient undesirable behavior to acquire a penalty count status.

The sender penalty count status module 27 determines if the incoming e-mail message 11 exhibits undesirable activity, at decision block 43, and if so, the sender penalty count status module 27 creates a new listing for the sender identifier 21 in a penalty count table 50 (shown in Fig. 3) with an appropriate sender penalty count value, at step 45. Or, if a listing already exists for the sender identifier 21, the corresponding sender penalty count value is updated. For normal e-mail activity in which no undesirable activity is detected, no change is made to the sender entry in the penalty count table 50. For a sender not listed in the penalty count table 50, the penalty count value is taken to be zero. In one embodiment of the present method, the activity of the sender of the e-mail message 11 is monitored after the TCP connection, established at step 31, has been terminated. Undesirable activity related to the current e-mail message 11 is noted, and the appropriate penalty count assessed. Operation of

the mail transfer agent 10 then proceeds to step 31, at which the next TCP connection is established.

Derivation of Penalty Count Status

The function of the penalty count table 50 can be described with reference to the data structure diagram in Fig. 3. The penalty count table 50 includes a plurality of records, represented by records 51, 53, 55, and 59. A semaphore feature 57, or a similar lock/unlock facility, is associated with each table entry for the purpose of synchronization when, for example, two or more processors are used to process the incoming e-mail message 11. The structure of the records 53, 55, and 59 are similar to that of the record 51 which includes a sender identifier (SID) value 61, a cumulative penalty count (PNCT) value 63, and a timestamp (TS) value 65.

In a preferred embodiment, the process of looking up sender identifier values can be optimized by using keys computed from the SIDs to hash the penalty count table 50. In a table with 256 entries, for example, the least significant byte of a corresponding sender identifier value can be used as the hash key. Rehashing can be minimized by a configuration in which each entry of the hash table points to a linked list, where the records are stored in the link list.

The cumulative penalty count value 63 is a time-dependent parameter which is updated in accordance with the behavior of the sender identified by the corresponding SID value 61. The timestamp value 65 records the time T_{TS} at which the cumulative penalty count value 63 was most recently calculated. The timestamp value 65 also provides for determining when the record 51 becomes out-of-date and should be removed from the penalty count table 50. In a preferred embodiment, the record 51 is removed after a retention period τ_{PCNT} of approximately 2^{19} seconds (i.e., about six days). The cumulative penalty count (PCNT) value 63 is preferably derived using the equation,

$$PCNT = \min(\kappa, \alpha + \varphi), \quad (1)$$

where κ is a pre-established maximum value for PCNT, α is an activity penalty count charged to the sender for the current undesirable activity, and φ is a previous penalty count value which was determined from the recent past history, if any, of undesirable activity produced by the sender. In a preferred embodiment, κ is set to
 5 128.

The value for the activity penalty count α can be an integer specified by the system administrator, and may have a different value for different types of undesirable activities. For example, an activity penalty count of six may be assessed for sending a large number of e-mails exceeding a pre-established maximum quantity,
 10 an activity penalty count of four may be assessed for sending one or more e-mails exceeding a pre-established cumulative file size, and another activity penalty count may be assessed for incurring a TCP connection time exceeding a predetermined threshold. The activity penalty count is additive such that a sender can be assessed an activity penalty count of ten for exceeding both the maximum quantity and file size,
 15 for example.

The process of deriving the cumulative penalty count value 63 begins with the occurrence of an initial undesirable activity, for which a first activity penalty count of α_1 is charged to the sender identified by the SID value 61. As described above, the timestamp value 65 records the time (denoted as T_1 in the following example) of the
 20 occurrence of the current (i.e., the first) undesirable activity. Accordingly,

$$PCNT(T_1) = \min(\kappa, \alpha_1) = \alpha_1, \quad (2)$$

since $\varphi = 0$ where there has occurred only the first undesirable activity.

If the sender identified by the SID value 61 produces a subsequent (i.e., the second) undesirable activity at a time T_2 , a second activity penalty count α_2 is assigned. If T_2 lies within the retention period τ_{PCNT} , the previous cumulative penalty
 25 count (PCNT) 63 is updated to the value $PCNT(T_2)$ using the expression,

$$PCNT(T_2) = \min (\kappa, \alpha_2 + \varphi_2), \quad (3)$$

where,

$$\varphi_2 = \left[PCNT(T_1) \cdot \left(1 - \frac{T_2 - T_1}{\tau_{PCNT}} \right) \right], \quad (4)$$

to give,

$$PCNT(T_2) = \alpha_2 + \alpha_1 \left(1 - \frac{T_2 - T_1}{\tau_{PCNT}} \right). \quad (5)$$

Note that, if T_2 occurs after the retention period τ_{PCNT} following T_1 , then $\varphi_2 = 0$.

In general, the n^{th} cumulative penalty count 63, updated at time T_n can be
 5 determined using the expression,

$$PCNT(T_n) = \min (\kappa, \alpha_{n-1} + \varphi_n), \quad (6)$$

where,

$$\varphi_n = \left[PCNT(T_{n-1}) \cdot \left(1 - \frac{T_n - T_{n-1}}{\tau_{PCNT}} \right) \right]. \quad (7)$$

Determination of PCFM State

In a preferred embodiment, the penalty count filter module 20 operates in one of at least three states: a 'normal' state, a 'selective-rejection' state, and a 'random-rejection' state. When the penalty count filter module 20 is operating in the 'normal' state, the mail transfer agent 10 accepts all valid incoming e-mail message 11 from any senders for transmittal to the intended addressees. As the system resources required to handle incoming e-mail volume increase and a greater demand is placed on the mail transfer agent 10, operation of the penalty count filter module 20 will change from the 'normal' state to either the 'selective rejection' state or the 'random rejection' state. Subsequently, if the e-mail volume returns to normal levels, operation of the penalty count filter module 20 eventually reverts to 'normal.'

If initially in the 'normal' state, operation of the penalty count filter module 20 changes from the 'normal' state to the 'selective-rejection' state if the system resources have increased beyond a first pre-established, 'selective-rejection' watermark. The selective-rejection watermark is reached, for example, when the disk space of the incoming message queue exceeds a predetermined disk-space threshold, or when the number of concurrent TCP connections exceeds a predetermined connection number. In alternative embodiments, other such criteria can be used, as specified by the system administrator, to define additional watermarks and operational states.

10 If initially operating in the 'random-rejection' state, the penalty count filter module 20 will remain in the 'random-rejection' state for at least a period of time denoted as a time-to-stay (TTS) interval τ_{TTS} . After the time-to-stay interval τ_{TTS} has passed, the penalty count filter module 20 may revert to either 'selective-rejection' operation or 'normal' operation, depending on the system resource usage status file
15 25. Alternatively, if initially operating in the 'selective-rejection' state, the penalty count filter module 20 will continue in the 'selective-rejection' operation for the time-to-stay interval τ_{TTS} and then revert to the 'normal' state, except when e-mail volume increases and operation of the penalty count filter module 20 is changed to the 'random-rejection' state. In a preferred embodiment, the time-to-stay interval τ_{TTS} is
20 approximately ten minutes.

The application of the PCFM state 29 to the process of managing the incoming e-mail message 11 can be explained with reference to the flow diagram of Fig. 4, which provides a more detailed description of step 33 in Fig. 2. From step 31, the system resource usage status 25 is determined, at step 71. A query is made as to
25 whether the penalty count filter module 20 is in the selective-rejection state, at decision block 73. If the response is 'yes,' operation proceeds to decision block 85. If the response is 'no' in decision block 73, a subsequent query is made as to whether

selective-rejection watermark has been exceeded, at decision block 93, the time-to-stay is updated, at step 97, the penalty count filter module 20 remains in the selective-rejection state, at step 99, and operation proceeds to step 35.

If the time-to-stay has not expired, at decision block 91, a query is made as to whether the time-to-check has expired, at decision block 101. If the time-to-check has not expired, the penalty count filter module 20 remains in the selective-rejection state, at step 99, and operation proceeds to step 35. If the time-to-check has expired, at decision block 101, a query is made as to whether the system resource has exceeded the selective-rejection watermark, at decision block 103. If the selective-rejection watermark has been exceeded, a resource usage factor f (defined below) is increased, at step 105, and the time-to-check is updated, at step 109. In one preferred embodiment, the resource usage factor f is doubled when the system resource exceeds the selective-rejection watermark. If the selective-rejection watermark has not been exceeded, at decision block 103, the resource usage factor f is decreased, at step 107, and the time-to-check is updated, at step 109. In another preferred embodiment, the resource usage factor f is decreased by a factor of two if the selective-rejection watermark has not been exceeded. After the time-to-check has been updated, at step 109, the penalty count filter module 20 remains in the selective-rejection state, at step 111, and operation proceeds to step 35.

If the penalty count filter module 20 is determined to be in the random-rejection state, at decision block 75 (above), operation proceeds to decision block 115 at which a query is made as to whether the time-to-stay has expired. If the time-to-stay has expired, a query is made as to whether the system resource has exceeded the random-rejection watermark, at decision block 117. If the random-rejection watermark has not been exceeded, the time-to-stay is initialized to the time-to-stay interval τ_{TTS} , the resource usage factor f is set to an initial value, and the time-to-check is initialized to the time-to-check interval τ_{CHK} , at step 119. Subsequently, the

penalty count filter module 20 changes to the selective-rejection state, at step 121, and operation proceeds to step 35. If the random-rejection watermark has been exceeded, at decision block 117, the time-to-stay is updated, at step 123, the penalty count filter module 20 remains in the random-rejection state, at step 125, and operation proceeds to step 35.

If the time-to-stay has not expired, at decision block 115, a query is made as to whether the time-to-check has expired, at decision block 127. If the time-to-check has not expired, the penalty count filter module 20 remains in the random-rejection state, at step 125, and operation proceeds to step 35. If the time-to-check has expired, at decision block 127, a query is made as to whether the system resource has exceeded the random-rejection watermark, at decision block 129. If the random-rejection watermark has been exceeded, the rejection factor R_f is increased, at step 131, and the time-to-check is updated, at step 135. In a preferred embodiment, the rejection factor R_f is doubled when the system resource exceeds the random-rejection watermark. If the selective-rejection watermark has not been exceeded, at decision block 129, the rejection factor R_f is decreased, at step 133, and the time-to-check is updated, at step 135. In one preferred embodiment, the rejection factor R_f is decreased by a factor of two if the random-rejection watermark has not been exceeded. After the time-to-check has been updated, at step 135, the penalty count filter module 20 remains in the random-rejection state, at step 137, and operation proceeds to step 35.

Processing E-mail

The operation of processing e-mail in accordance with sender status and the penalty count filter module state, at step 39 of Fig. 2, is shown in greater detail in the flow diagram of Fig. 5, in which a query is made, at decision block 201, as to whether the penalty count filter module 20 is in the selective-rejection state. If the response is 'yes,' a query is made as to whether the current value of the cumulative penalty count

φ is greater than zero, at decision block 203. If the response is 'no,' the incoming e-mail message 11 is accepted, at step 171, and sent to the addressee as forwarded e-mail 17. If the response is 'yes,' in decision block 203, a random number R_h is generated, at step 205, where $1 \leq R_h \leq \kappa$.

- 5 A query is then made, at decision block 207, as to whether the random number R_h is greater than the product of the resource usage factor f and the current value of the cumulative penalty count φ . If the response is 'yes,' the incoming e-mail message 11 is accepted, at step 171, and forwarded to the addressee. Operation returns to step 41 where the behavior trace table 160 is updated. If the response is
- 10 'no,' at decision block 207, the incoming e-mail message 11 is rejected, at step 209, the penalty count filter module 20 returns the transient negative completion reply 19 to the sender, and operation returns to step 41.

- If the response is 'no,' at decision block 201, a query is made as to whether the penalty count filter module 20 is in the random-rejection state, at decision block 211.
- 15 If the response is 'no,' the incoming e-mail message 11 is accepted, at step 171, and sent to the addressee as forwarded e-mail 17. If the response is 'yes,' at decision block 211, a query is made as to whether the current value of the cumulative penalty count φ is greater than zero, at decision block 213. If the response is 'yes,' at decision block 213, the incoming e-mail message 11 is rejected, at step 219, and the
- 20 penalty count filter module 20 returns the transient negative completion reply 19 to the sender. If the response is 'no,' at decision block 213, a random number R_g is generated, at step 215, where $1 \leq R_g \leq \kappa$.

- A query is made, at decision block 217, as to whether the random number R_g is greater than the rejection factor R_f . If the response is 'yes,' the incoming e-mail
- 25 message 11 is accepted, at step 171, and forwarded to the addressee. If the response is 'no,' at decision block 217, the incoming e-mail message 11 is rejected, at step 219, the penalty count filter module 20 returns the transient negative completion reply 19

to the sender, and operation returns to step 41 where the behavior trace table 160 is updated.

As described above, the resource usage factor f is used at step 207 to determine whether a particular incoming e-mail message 11 is to be randomly
 5 returned to the sender or transmitted to the intended addressee. The resource usage factor f is derived from the average cumulative penalty count \bar{P} , where

$$\bar{P} = \frac{\sum_{i=1}^m PCNT_i}{m}, \quad (8)$$

and where $PCNT_i$ is the penalty count assessed to the i^{th} sender identifier, of the m sender identifiers listed in the penalty count table 50. The resource usage factor f is given by the equation,

$$f = \frac{\kappa}{2\bar{P}}. \quad (9)$$

10 As described above, the penalty count filter module 20 then generates the random number R_h , which is compared to the product of the resource usage factor f and the current value of the cumulative penalty count ϕ . If $R_h > f \cdot \phi$, the incoming e-mail message 11 is accepted by the penalty count filter module 20; if $R_h \leq f \cdot \phi$, the incoming e-mail message 11 is rejected and the penalty count filter module 20 issues
 15 the transient negative completion reply 19 to the corresponding sender.

Behavior Trace Table

In a preferred embodiment, the behavior trace table 160 is included in a workstation 150 which is connected with the mail transfer agent 10 by means of a communication network 151, such as a LAN or WAN, as shown in Fig. 6.
 20 Information included in the behavior trace table 160, which tabulates e-mail activities of the sender, is used by the sender penalty count status 27 to update the cumulative penalty count ϕ of an e-mail sender in the penalty count table 50.

The behavior trace table 160 includes a plurality of records, or behavior values, represented by records 161, 163, 165 and 169, shown in Fig. 7. A semaphore feature 167, or a similar lock/unlock facility, is associated with each table entry. The structure of the records 163, 165, and 169 are similar to that of the record 161 which includes a sender identifier (SID) value 171, a cumulative e-mail count (NN) 173 of the sender e-mails, a total e-mail size (SZ) 175 of the e-mail files, a total TCP connection utilization time (UT) 177, and a timestamp (TS) value 179 which is a record of the most recent e-mail received from the sender identified by the SID value 171.

The process of looking up sender identifier values can be optimized by using keys computed from the SIDs to hash the behavior trace table 160, using a method similar to that for looking up values in the penalty count table 50. Rehashing can be minimized in the behavior trace table 160 by having each entry of the hash table point to a linked list, where the records are stored in the link list. Each record 161, 163, 167, and 169 in the behavior trace table 160 has a retention period of τ_{BHT} , after which the out-of-date record is removed when the corresponding table entry is visited and the record is traversed. In a preferred embodiment, the retention period τ_{BHT} is approximately five seconds.

When the incoming e-mail message 11 corresponds to the sender identified by the sender identifier value 171, the cumulative e-mail count 173, the cumulative total size 175, and the cumulative total TCP connection utilization time 177 are updated using current and previous values, and where the previous values are reduced by a decay factor d . The updated cumulative e-mail value (NN) is given by the equation,

$$NN = 1 + d \cdot NN(T_{TS}), \quad (10)$$

where $NN(T_{TS})$ is the prior or most recent previous e-mail count value, previously obtained at the timestamp time (T_{TS}).

The updated cumulative total size (SZ) is given by the equation,

$$SZ = SZ(T_{UD}) + d \cdot SZ(T_{TS}), \quad (11)$$

where $SZ(T_{UD})$ is the additional e-mail size obtained at the time of updating the record and $SZ(T_{TS})$ is the prior or most recent previous e-mail size. Similarly, the updated cumulative connection utilization time (UT) is given by the equation,

$$UT = UT(T_{UD}) + d \cdot UT(T_{TS}), \quad (12)$$

where $UT(T_{UD})$ is the connection utilization time determined at the time of updating the record and $UT(T_{TS})$ is the prior or most recent previous connection utilization time. After the cumulative e-mail count 173, the cumulative total size 175, and the cumulative total TCP connection utilization time 177 have been updated, the timestamp (T_{TS}) is reset to the most recent time of update. The decay factor d in equations 10 to 12 is given by,

$$d = 1 - \min \left(1, \frac{T_{TC} - T_{TS}}{\tau_{BHT}} \right). \quad (13)$$

10 If any of the cumulative e-mail count 173, the cumulative total size 175, or the cumulative total TCP connection utilization time 177 exceeds a predefined threshold or watermark, as discussed above, the sender corresponding to the sender identifier value 171 is assessed an appropriate penalty count. For incoming e-mail message 11 sent by a sender not currently listed in the behavior trace table 160, a new entry is
15 created with initial values of one, $SZ(T_{TS})$, and $UT(T_{TS})$ entered into the behavior trace table 160 for the cumulative e-mail count 173, the cumulative total size 175, and the cumulative total TCP connection utilization time 177, respectively.

In yet another embodiment, the mail transfer agent 10 includes a merged penalty count table 180, shown in Fig. 8, which includes a plurality of records having
20 entries similar to that shown for a record 181. The record 181 results from merging the data in the record 51 and the record 161, and includes a sender identifier (SID) value 183, a cumulative penalty count (PCNT) value 185, a cumulative e-mail count (NN) 187 of the sender e-mails, a total e-mail size (SZ) 189 of the e-mail files, a total

TCP connection utilization time (UT) 191, and a timestamp (TS) value 193. As can be appreciated by one skilled in the relevant art, a behavior trace table 160 is not required in the workstation 150 if the mail transfer agent 10 includes the merged penalty count table 180 in place of the penalty count table 50.

5 In an alternative embodiment, shown in the flow diagram of Fig. 9, the incoming e-mail 11 is processed in accordance with sender status. The mail transfer agent 10 completes a TCP connection, at step 231, and receives the incoming e-mail message 11. The identity of the sender of the e-mail message 11 is determined, at step 233. The accept/reject filter 23 determines, at step 235, either that the incoming
10 e-mail message 11 is to be transmitted to the addressee or that the incoming e-mail message 11 is to be rejected. The determination of accepting or rejecting the e-mail message 11 is made based on the identity of the sender. When the current incoming e-mail message 11 has been processed, the behavior trace table 160 is updated, at step 239, and the sender penalty count status module 27 determines if the incoming e-mail
15 message 11 exhibits undesirable activity, at decision block 241. If so, the sender penalty count status module 27 creates a new listing for the sender identifier 21 in the penalty count table 50 with an appropriate sender penalty count value, at step 243. Or, if a listing already exists for the sender identifier 21, the corresponding sender penalty count value is updated. For normal e-mail activity in which no undesirable
20 activity is detected, no change is made to the sender entry in the penalty count table 50. For a sender not listed in the penalty count table 50, the penalty count value is taken to be zero. Operation of the mail transfer agent 10 then proceeds to step 231, at which the next TCP connection is established.

25 While the invention has been described with reference to particular embodiments, it will be understood that the present invention is by no means limited to the particular constructions and methods herein disclosed and/or shown in the

drawings, but also comprises any modifications or equivalents within the scope of the claims.

What is claimed is:

CLAIMS

1 1. A method suitable for use in a communication device for determining the
2 disposition of incoming e-mail from a sender, said method including the step of
3 establishing the identity of the sender to provide a sender identifier (61) and
4 characterized in that said method further comprises the steps of:

5 determining a cumulative penalty count value (63) associated with said sender
6 identifier;

7 retrieving a system resource usage status (71) associated with the
8 communication device; and

9 processing the incoming e-mail on the basis of said cumulative penalty count
10 value and said system resource usage status.

1 2. The method of claim 1 wherein said step (35) of establishing the identity of
2 the sender comprises the step of ascertaining an IP address for the sender.

1 3. The method of claim 1 wherein said step (35) of establishing the identity of
2 the sender comprises the step of associating the sender with a peer IP address of the
3 sender TCP connection.

1 4. The method of claim 1 wherein said step of determining a cumulative penalty
2 count value comprises the step of assessing a penalty count value to said sender
3 identifier for an undesirable activity associated with the sender.

1 5. The method of claim 4 wherein said cumulative penalty count value comprises
2 an activity penalty count charged to the sender for current undesirable sender activity
3 and a time-dependent penalty count determined from previous undesirable sender
4 activity.

1 6. The method of claim 5 wherein said time-dependent penalty count comprises a
2 zero value subsequent to a pre-established retention period.

1 7. The method of claim 5 wherein said time-dependent penalty count comprises a
2 prior activity penalty count value reduced by a decay factor.

1 8. The method of claim 4 wherein said undesirable activity comprises a member
2 of the group consisting of: sending a large number of e-mails, sending e-mails of
3 relatively large sizes, using a relatively large amount of TCP connection time, and
4 causing a TCP timeout.

1 9. The method of claim 1 wherein said system resource usage status is a function
2 of a member of the group consisting of: the number of concurrent TCP connections
3 being maintained, the number of e-mail files in an incoming message queue, and the
4 amount of disk space being utilized for an incoming message queue.

1 10. The method of claim 1 wherein said step of processing the incoming e-mail
2 comprises the step of assigning an operating state to the communication device, said
3 operating state being a function of said system resource usage status.

1 11. The method of claim 10 wherein said operating state is a member of the group
2 consisting of: a normal operating state, a selective-rejection operating state, and a
3 random-rejection operating state.

1 12. The method of claim 11 wherein, for said selective-rejection state, if said
2 cumulative penalty count value has a zero value, said step of processing the incoming
3 e-mail comprises the step of accepting the incoming e-mail.

1 13. The method of claim 11 wherein, for said selective-rejection state, if said
2 cumulative penalty count value has a nonzero value, said step of processing the
3 incoming e-mail comprises the steps of:

4 specifying a rejection factor;
5 generating a random number; and
6 randomly rejecting the incoming e-mail on the basis of said rejection factor and
7 said random number.

1 14. The method of claim 13 wherein said step of randomly rejecting comprises the
2 step of accepting the incoming e-mail if said random number is greater than said
3 rejection factor and rejecting the incoming e-mail if said random number is not
4 greater than said rejection factor.

1 15. The method of claim 13 wherein said rejection factor is increased if said
2 system resource usage status increases and said rejection factor is decreased if said
3 system resource usage status decreases.

1 16. The method of claim 11 wherein, for said random-rejection state, if said
2 cumulative penalty count value has a nonzero value, said step of processing the
3 incoming e-mail comprises the step of rejecting the incoming e-mail.

1 17. The method of claim 11 wherein, for said random-rejection state, if said
2 cumulative penalty count value has a zero value, said step of processing the incoming
3 e-mail comprises the steps of:

4 deriving a resource usage factor;
5 generating a random number; and

6 randomly rejecting the incoming e-mail on the basis of said resource usage
7 factor, said random number, and said cumulative penalty count value.

1 18. The method of claim 17 wherein said step of randomly rejecting comprises the
2 step of accepting the incoming e-mail if said random number is greater than a product
3 of said resource usage factor and said cumulative penalty count value, and rejecting
4 the incoming e-mail if said random number is not greater than said product of said
5 resource usage factor and said cumulative penalty count value.

1 19. The method of claim 18 wherein said resource usage factor is increased if said
2 system resource usage status increases and said resource usage factor is decreased if
3 said system resource usage status decreases.

1 20. A communication device for determining the disposition of incoming e-mail
2 from a sender, said device including means for identifying the sender (21),
3 characterized in that said device further comprises:

4 a penalty count filter module (20) having
5 means for assigning a penalty count (27) to the sender, said
6 penalty count being a function of undesirable activity
7 associated with the sender;
8 means for determining a resource usage status (25) for said
9 communication device in receiving e-mail;
10 means for specifying an operating state (29) for said penalty
11 count filter module, said operating state being a function of
12 said resource usage value; and
13 an accept/reject filter (23) for disposing of the incoming e-mail
14 on the basis of said sender penalty count and said operating
15 state.

1 21. The device of claim 20 wherein said means for identifying the sender (21)
2 includes means for obtaining at least one of a Domain Name Service verification and
3 a peer IP address of the sender TCP connection.

1 22. The device of claim 20 wherein said undesirable activity comprises a member
2 of the group consisting of: sending a large number of e-mails, sending e-mails of
3 relatively large sizes, using a relatively large amount of TCP connection time, and
4 causing a TCP timeout.

1 23. The device of claim 20 wherein said system resource usage status (25) is a
2 function of a member of the group consisting of: the number of concurrent TCP
3 connections being maintained, the number of e-mail files in an incoming message
4 queue, and the amount of disk space being utilized for an incoming message queue.

1 24. A communication device for determining the disposition of incoming e-mail
2 from a sender, characterized in that said device comprises:

3 a sender penalty count data structure (27) for storing a current penalty count
4 value associated with the sender;

5 a system resource usage status file (25) for storing a current usage status value
6 for device e-mail processing resources; and

7 an accept/reject filter (23) for disposing of the incoming e-mail on the basis of
8 said penalty count value and said usage status.

1 25. The device of claim 24 wherein said sender penalty count data structure (27)
2 includes an entry comprising a member of the group consisting of: a sender
3 identification value, a cumulative penalty count value, a cumulative e-mail count, a
4 total e-mail size, a total TCP connection time, and a timestamp value.

1 26. A method suitable for use in a communication device for determining the
2 disposition of incoming e-mail from a sender, said method including the step of
3 identifying the e-mail sender by determining a sender IP address, characterized in that
4 said method further comprises the steps of:

5 obtaining a previous sender penalty count value calculated for said sender IP
6 address; and

7 accepting or rejecting the incoming e-mail based on said sender penalty count
8 value.

1 27. The method of claim 26 further comprising the steps of:

2 maintaining a behavior trace table entry for the e-mail sender; and

3 determining said previous sender penalty count from said behavior trace table.

1 28. The method of claim 27 further comprising the step of updating sender
2 behavior values in said trace table entry in response to receipt of a sender e-mail.

1 29. The method of claim 28 wherein said sender behavior values include a
2 member of the group consisting of: the number of e-mails, the total size of e-mails,
3 and the total time of TCP connection time.

1 30. The method of claim 28 wherein said step of updating sender behavior values
2 comprises the steps of:

3 reducing said behavior trace table value by a decay factor; and

4 adding a current behavior trace table value to said corresponding reduced
5 behavior trace table value.

1 31. The method of claim 30 wherein said decay factor is a function of the time
2 interval between the last two updates of said behavior trace table entry and a pre-
3 established retention period.

1 32. The method of claim 26 wherein said sender penalty count value is determined
2 from undesirable sender activity occurring over a pre-established retention period.

1 33. The method of claim 32 wherein said undesirable activity comprises a member
2 of the group consisting of: sending a large number of e-mails, sending e-mails of
3 relatively large sizes, using a relatively large amount of TCP connection time, and
4 causing a TCP timeout.

1 34. The method of claim 26 further comprising the step of updating said sender
2 penalty count value.

1 35. The method of claim 34 wherein said step of updating said sender penalty
2 count value comprises the steps of:

3 reducing said previous sender penalty count value by a decay factor to yield a
4 reduced sender penalty count value, said decay factor being a function of
5 said pre-established retention period; and
6 adding an activity penalty count value to said reduced sender penalty count
7 value to yield an updated sender penalty count value, said activity penalty
8 count value calculated as a function of current sender e-mail activities.

- 1 36. The method of claim 35 wherein said decay factor is further a function of the
- 2 time interval between calculation of said previous sender penalty count value and
- 3 calculation of said activity penalty count value.


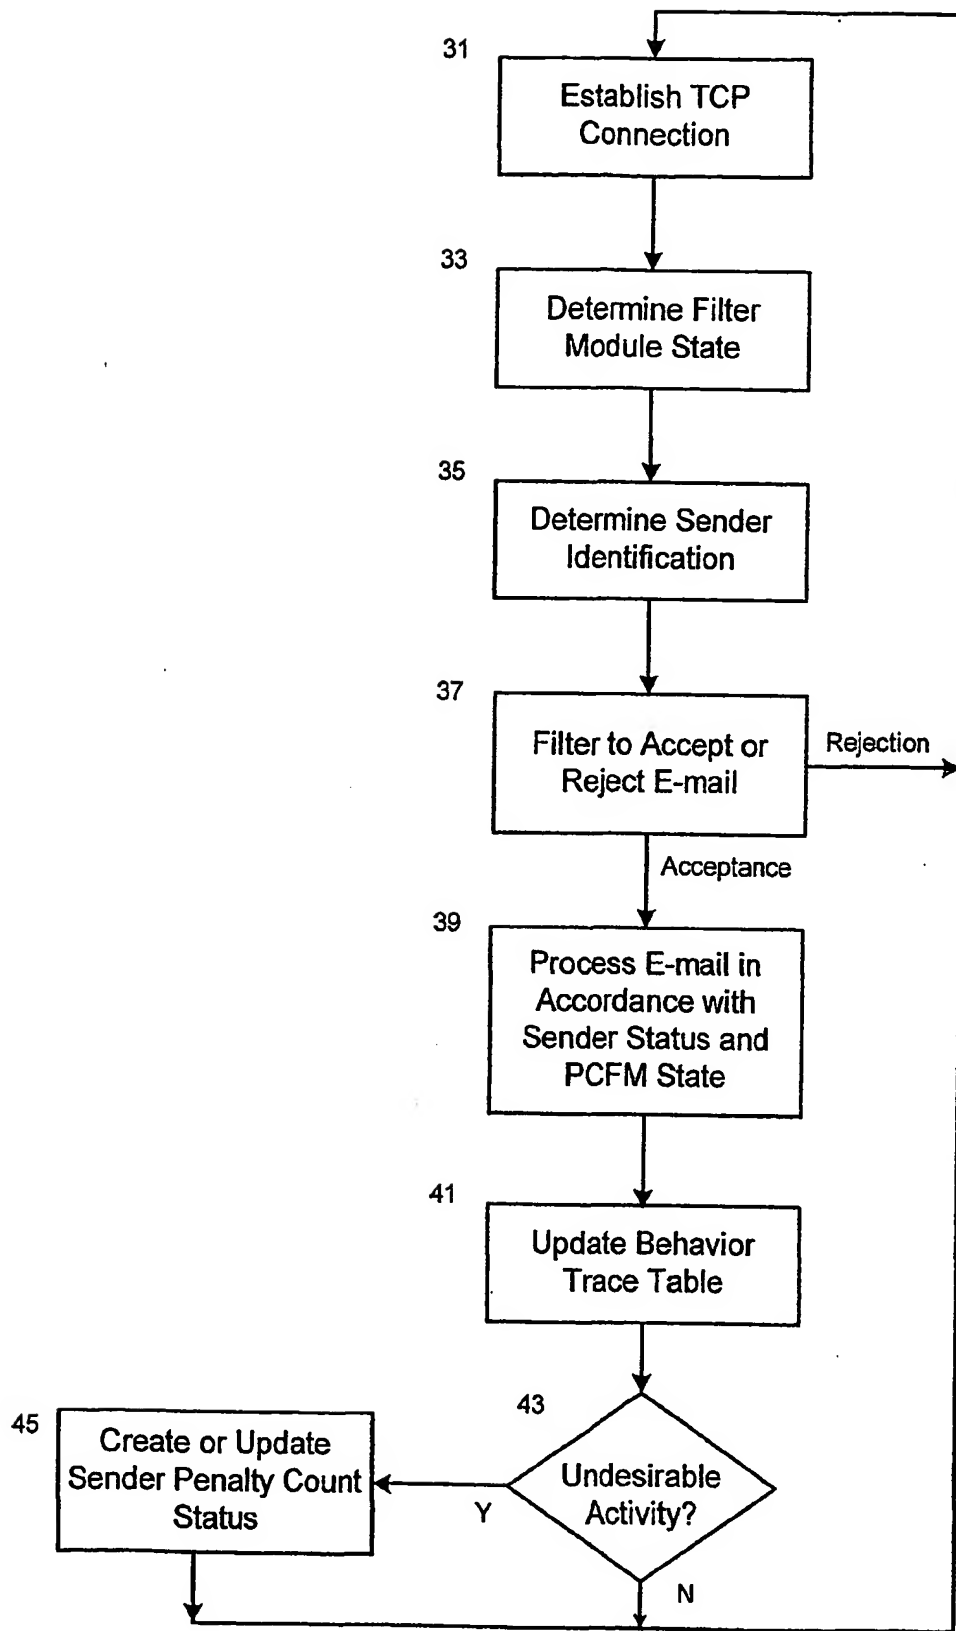


FIG. 1

**Fig. 2**


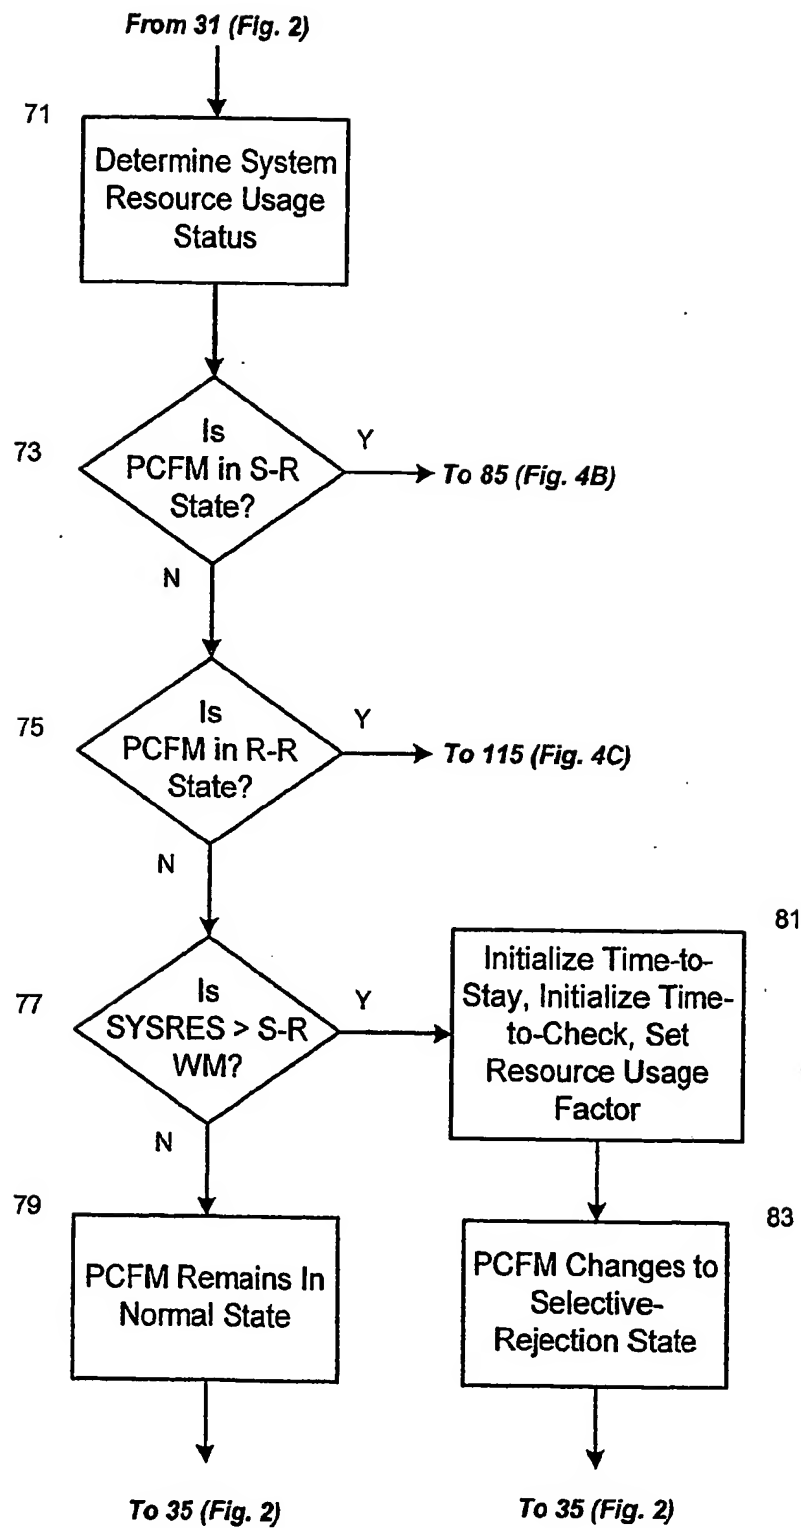


FIG. 3

**Fig. 4A**

5/10

From 73 (Fig. 4A)

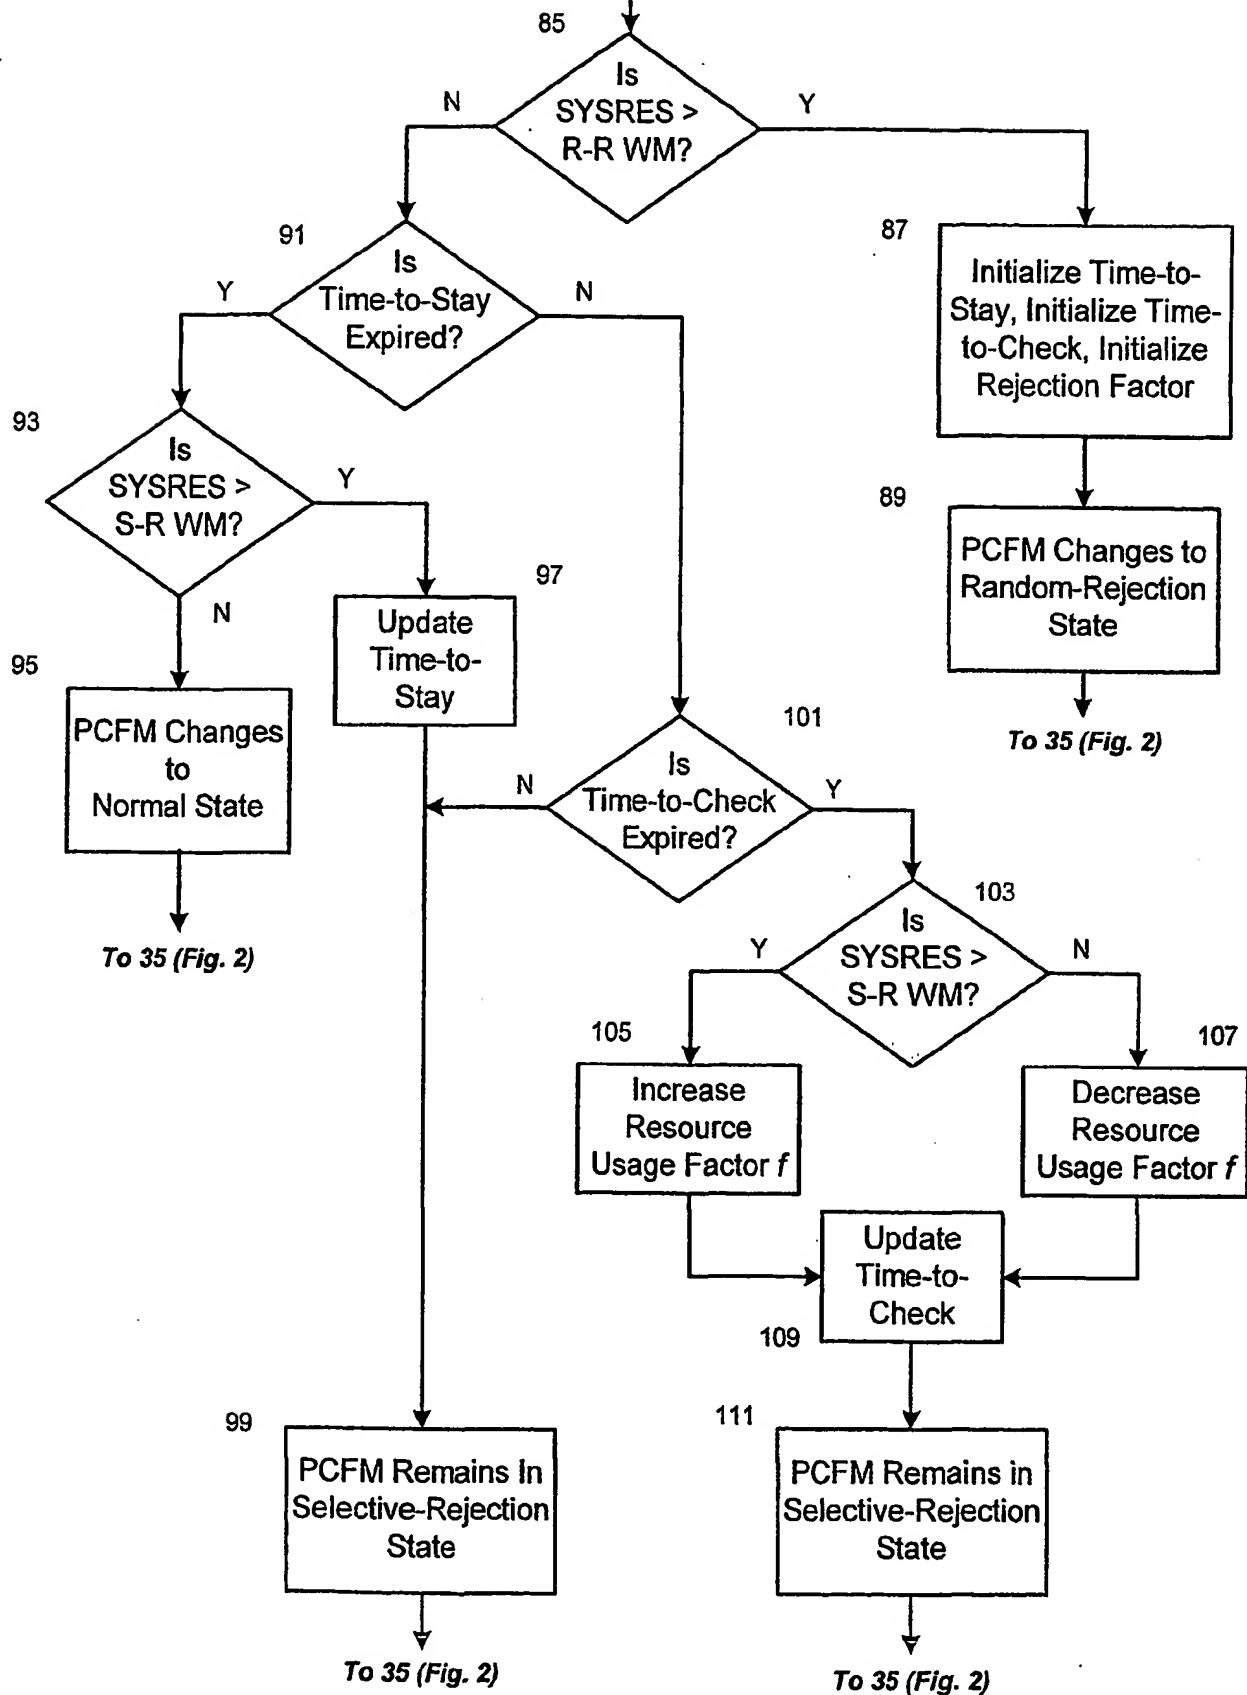


Fig. 4B

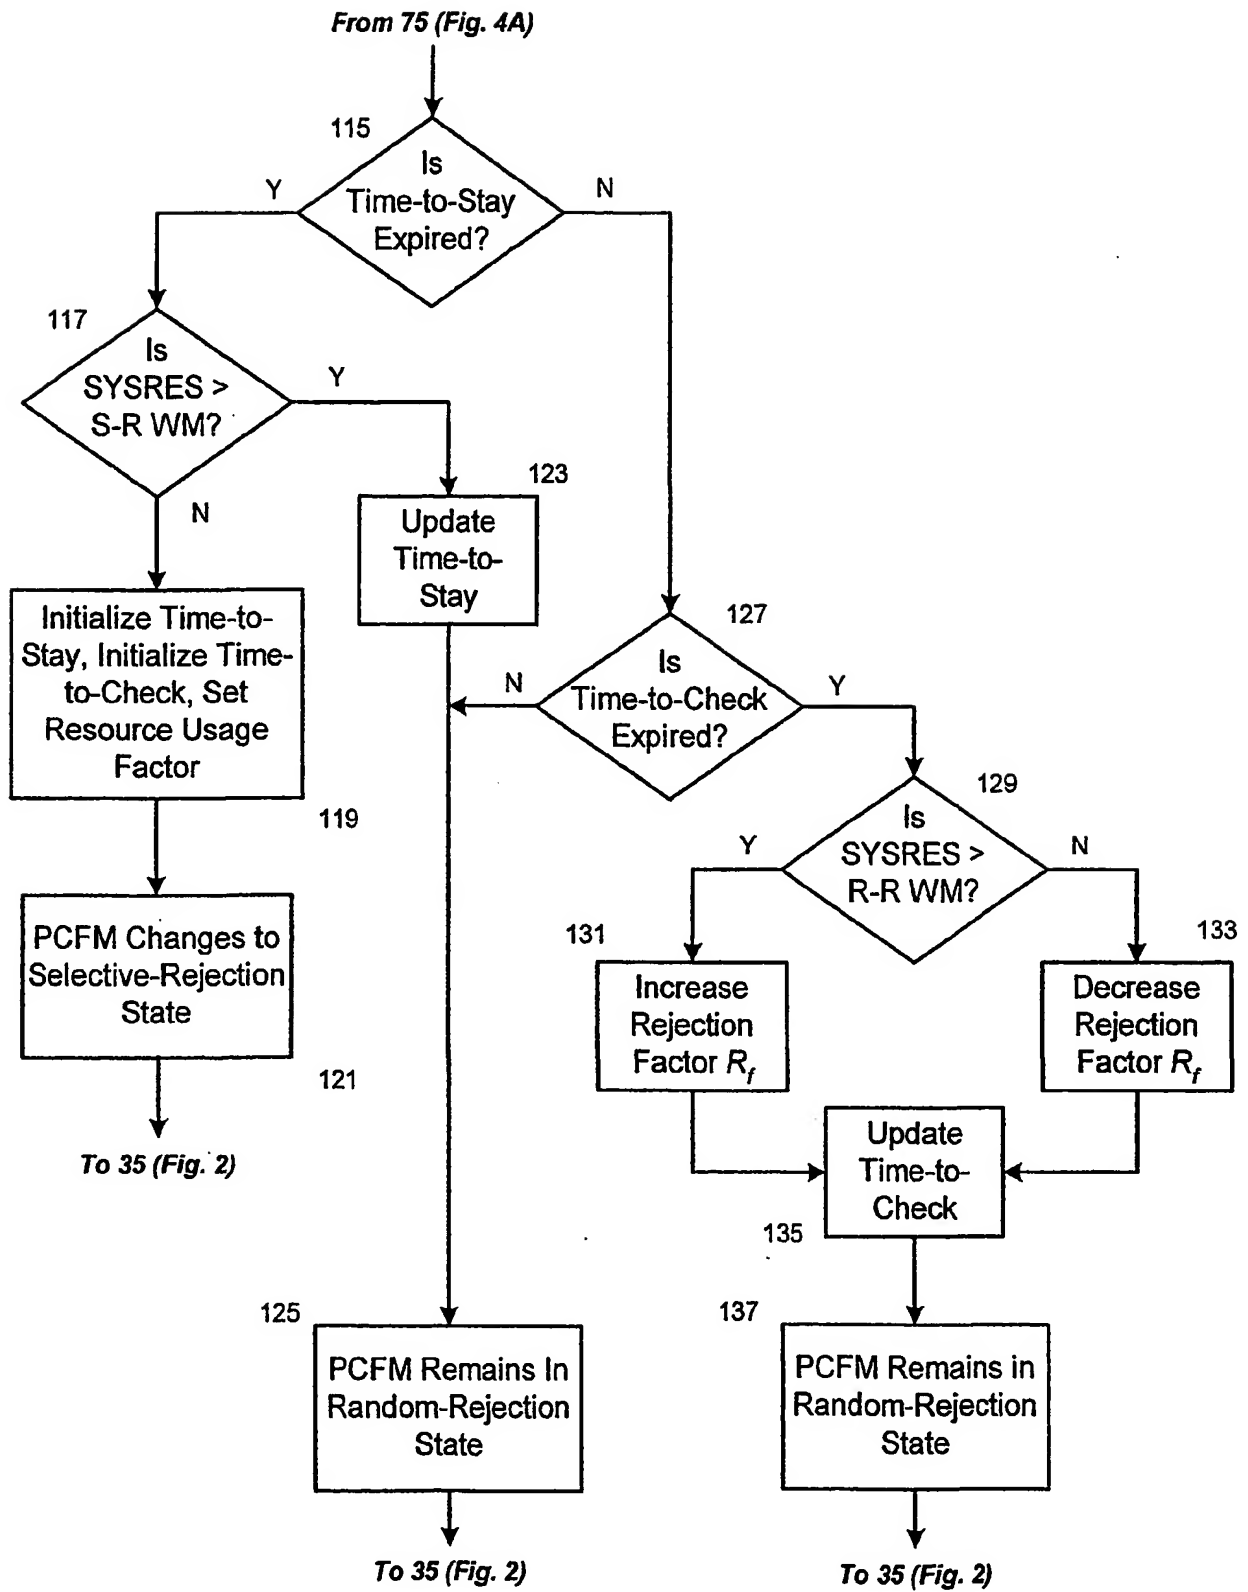



Fig. 4C

**Fig. 5**




FIG. 6




FIG. 7


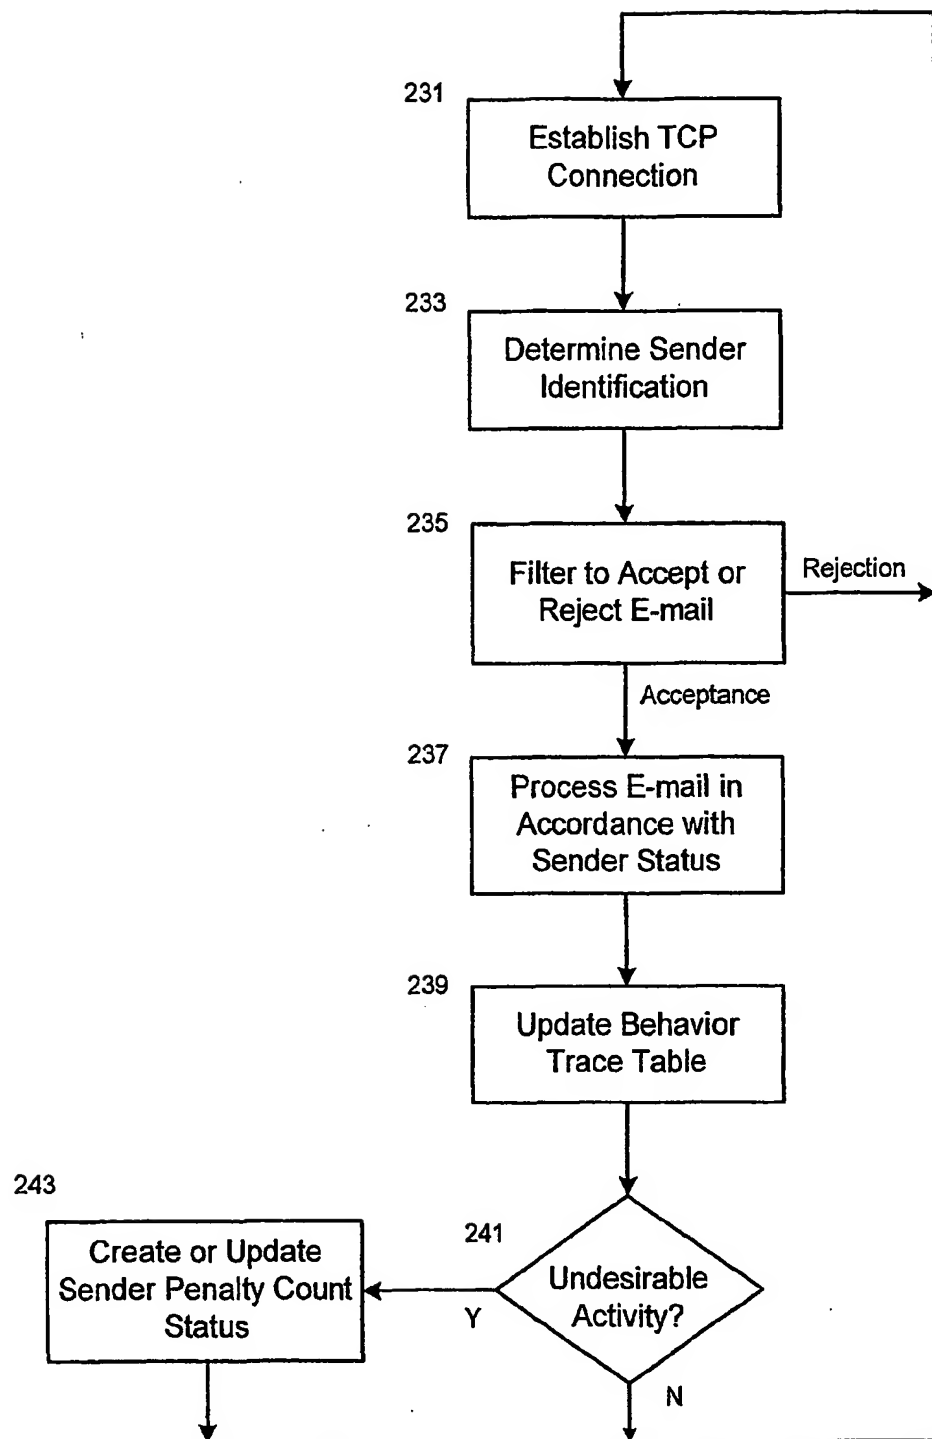


FIG. 8

**Fig. 9**

INTERNATIONAL SEARCH REPORT

International application No.

PCT/IB02/02507

A. CLASSIFICATION OF SUBJECT MATTER

IPC(7) : G06F 15/173, 17/00

US CL : 709/206, 229

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 709/206, 226, 229

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

USPAT

C. DOCUMENTS CONSIDERED TO BE RELEVANT

| Category * | Citation of document, with indication, where appropriate, of the relevant passages | Relevant to claim No. |
|------------|--|-----------------------|
| A | US 5,999,967 A (SUNDSTED) 07 December 1999 (07.12.99), abstract. | 1-36 |
| A,P | US 6,393,464 A (DIETERMAN) 21 May 2002 (21.05.2002), abstract | 1-36 |
| A,P | US 6,484,197 A (DONOHUE) 19 November 2002 (19.11.2002), abstract. | 1-36 |
| A | US 6,137,777 A (VAID et al.) 24 October 2000 (24.10.2000), abstract | 1-36 |
| A | US 5,315,586 A (CHARVILLAT) 24 May 1994 (24.05.1994), abstract | 1-36 |

☐ Further documents are listed in the continuation of Box C.

☐ See patent family annex.

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T"

later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X"

document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y"

document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&"

document member of the same patent family

Date of the actual completion of the international search

27 November 2002 (27.11.2002)

Date of mailing of the international search report

12 DEC 2002

Name and mailing address of the ISA/US

Commissioner of Patents and Trademarks

Box PCT

Washington, D.C. 20231

Facsimile No. (703)305-3230

Authorized officer

Dung Dinh

Telephone No. 305 9600